

A Policy Driven AI-Assisted PoW Framework

Trisha Chakraborty*, Shaswata Mitra†, Sudip Mittal‡, Maxwell Young§

Department of Computer Science & Engineering, Mississippi State University

*tc2006@msstate.edu, †sm3843@msstate.edu, ‡mittal@cse.msstate.edu, §myoung@cse.msstate.edu

Abstract—Proof of Work (PoW) based cyberdefense systems require incoming network requests to expend effort solving an arbitrary mathematical puzzle. Current state of the art is unable to differentiate between trustworthy and untrustworthy connections, requiring all to solve complex puzzles. In this paper, we introduce an Artificial Intelligence (AI)-assisted PoW framework that utilizes IP traffic based features to inform an ‘adaptive’ issuer which can then generate puzzles with varying hardness. The modular framework uses these capabilities to ensure that untrustworthy clients solve harder puzzles thereby incurring longer latency than authentic requests to receive a response from the server. Our preliminary findings reveal our approach effectively throttles untrustworthy traffic.

I. INTRODUCTION

A distributed denial-of-service (DDoS) attack is a malicious attempt to disrupt the normal traffic of a targeted server, service, or network by overwhelming the target or its surrounding infrastructure with a flood of Internet traffic. A possible defensive strategy is an effective Proof-of-Work (PoW) based system [3], [1], [2]. A PoW system works by requiring incoming network requests to expend effort solving an arbitrary mathematical puzzle to prevent anybody from attacking the system. In PoW based systems, a client has to commit some computation (CPU cycle, bandwidth, etc.) to solve a puzzle to prove her authenticity.

PoW systems generally consist of three parts: *issuer*, *solver*, and *verifier*. The issuer (also called a generator) issues the puzzle to the solver, which solves them and sends the solution to the verifier. In a simple networked client-server environment, the server contains the issuer/generator and the verifier components, and the client is the solver.

In this paper, we build an Artificial Intelligence (AI)-assisted PoW framework. We create an ‘adaptive’ issuer which can generate puzzles with varying hardness. The idea behind the system is to penalize untrustworthy connections by issuing them ‘hard’ puzzles and at the same time give ‘easy’ puzzles to trustworthy requests. *Critically, these challenges introduce latency in the environment for untrustworthy connections.* The distinction between trustworthy/untrustworthy clients can be made by using the incoming traffic specific features. In other words, an AI subsystem can compute a *reputation score* for an incoming request, that can guide the puzzle generator.

Our framework has two useful properties. First, each client pays a cost for utilizing the system, and this cost increases as the client’s reputation score worsens. Second, the amount of work inflicted by a puzzle is *adaptive* and can be tuned. The framework will ensure that clients with bad reputation scores incur longer latency to receive a response from the server than

authentic requests. This latency is beneficial to the network under attack to slow down the incoming malicious traffic.

By creating this AI-assisted PoW puzzling scheme with the above mentioned properties, we strengthen the defensive posture of organizations that use these strategies. Our framework is modular and each component can be customized. Components include: an AI model that generates a reputation score, a puzzle generator, a puzzle solver, a puzzle verifier, and a policy that maps a reputation score to puzzle hardness. Regarding the last component, a network administrator may specify a policy based on her specific security needs. Next, we describe our system architecture and evaluations. Our results demonstrate that the AI assisted PoW puzzle introduces latency in the network for clients with bad reputation scores.

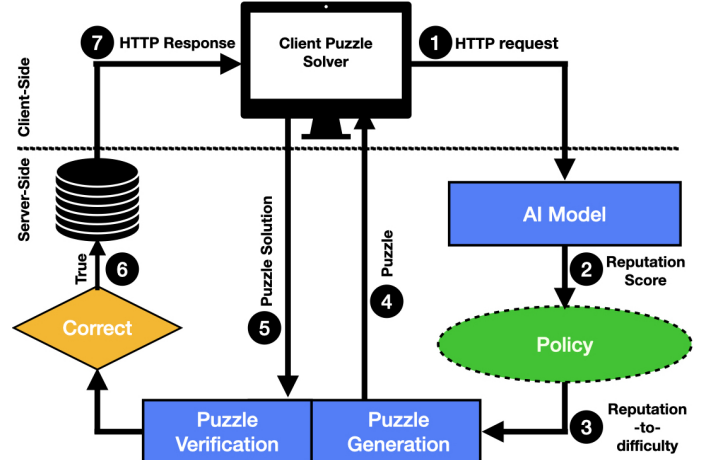


Fig. 1. An illustration of our framework. (1) The client issues an HTTP request to the server. (2) The AI model inspects the features of the request as input and produces a reputation score. (3) The policy module maps the reputation score to a difficulty. (4) The puzzle generation module generates a puzzle and sends it to the client to solve. (5) The puzzle verification module verifies the puzzle solution returned by the client. (6) On solving the puzzle correctly, the server is informed. (7) The server responds with the requested resource.

II. SYSTEM ARCHITECTURE

The architecture of our AI-assisted PoW framework is depicted in Figure 1. The components of this framework are described below.

1) *AI Model*: To create a proof of concept we use DABR [4] as our AI model. DABR is an euclidean distance-based technique that generates a reputation score for an IP address by learning from previously known malicious IP addresses and their attributes. The model generates a reputation score for an IP with an accuracy of 80%. The reputation score is normalized to a scale of 0 - 10, where a higher score represents

a more untrustworthy client. This reputation score serves as an input to the policy module as a variable mapped to puzzle hardness.

2) *Policy*: The policy module takes as an input a reputation score R (range $[0, 10]$). Intuitively, a high reputation score of 10, denoting an untrustworthy client should be allotted a harder puzzle than a relatively better reputation score of 9. Hence, proper tuning of the difficulty is desired for fine-grained reputation scores. A *policy* is a rule-based strategy for mapping the reputation score of a client to the appropriate puzzle difficulty. In Section III, we implement and analyze three such policies.

3) *Puzzle generation*: Our puzzle generation module issues a PoW puzzle. A puzzle is generated by collecting request related data, i.e., timestamp and unique seed (for mitigating pre-computation attacks), and a difficulty value as defined by the policy module, all of which is relayed back to the client.

4) *Puzzle solver*: The data received from the puzzle generation module are concatenated with the client's IP address to form a string that is not altered. To this, a 32-bit string is added, which the client modifies upon each hash function evaluation. The client performs evaluations on this input until it finds an output with a prefix of d zeros; we refer to this as a **d-difficult puzzle**. Our strategy is to assign higher difficulty of puzzles to clients with bad reputations scores.

5) *Puzzle Verification*: Puzzle verification is light weight block used to verify the clients solution and offer response if correct solution is returned.

III. EVALUATION

For our proof of concept implementation we utilize DAbR [4], which uses IP protocol based features to generate a reputation score. In this section, we discuss three policies used for a preliminary evaluation of our framework.

A. Policies 1 and 2: Linear mapping

For this evaluation, R take on values in the set $\{0, 1, \dots, 10\}$. The lower the reputation score, the more confidence we have that the client is trustworthy; conversely, the higher the reputation score, the more we suspect that the client is untrustworthy. Consequently, the difficulty of a puzzle assigned to a client increases with the client's reputation score. It takes 31 ms on average to solve a 1-difficult puzzle, and this time increases with difficulty.

For Policy 1, we map a 1-difficult puzzle to a client with a reputation score 0, a 2-difficult puzzle to a client with a reputation score of 1, and so on. Figure 2 shows that the latency increases with the increase in reputation score. However, the latency does not grow significantly as the reputation score increases. To address this, we evaluate Policy 2, where the easiest puzzle has difficulty 5. Thus, we map a 5-difficult puzzle to the client with reputation score 0, a 6-difficult puzzle to a client with a reputation score of 1, and so on. As a result, the latency increases significantly with higher reputation scores, delaying service for untrustworthy clients.




Fig. 2. An evaluation of our three implemented policies. The median of 30 trials is reported for each reputation score.

B. Policy 3: Error range mapping

In this policy, we consider the error ϵ from DAbR system [4]. Note that, given this error, the resulting IP reputation score might be higher or lower than the ground truth. Our Policy 3 attempts to correct for this in the following way. All reputation scores s_i are in the interval $[0, 10]$. For a score s_i , the difficulty value is a value chosen at random in the interval $[[d_i - \epsilon], [d_i + \epsilon]]$, where $d_i = \lceil s_i + 1 \rceil$. Figure 2 shows how the rate of increase in the latency for Policy 3 is between our two previous policies.

Summary. Our evaluation illustrates how the amount of latency can be tuned given different mappings between reputation scores and puzzle difficulty. Thus, our framework offers flexibility and can accommodate the specific security demands for a range of network settings.

IV. CONCLUSION

In this paper, we proposed and evaluated an AI-assisted PoW framework. It employs an AI model that uses IP-based features to output a reputation score. This score serves as an input to an ‘adaptive’ puzzle generator which creates puzzles of varying difficulty. In this way, the framework imposes higher latency on untrustworthy clients. We evaluated our proposed framework, showcasing how it can help defend organizations against networking threats.

Acknowledgements. This work was supported in part by NSF grant CNS-1816076.

REFERENCES

- [1] Diksha Gupta, Jared Saia, and Maxwell Young. Proof of work without all the work. *Proceedings of the 19th International Conference on Distributed Computing and Networking*, 2018.
- [2] Tien Le, Akshay Dua, and Wu-chang Feng. kapow plugins: protecting web applications using reputation-based proof-of-work. In *2nd Joint WICOW/AIRWeb Workshop on Web Quality*, pages 60–63, 2012.
- [3] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. <http://bitcoin.org/bitcoin.pdf>.
- [4] Arya Renjan, Karuna Pande Joshi, Sandeep Nair Narayanan, and Anupam Joshi. Dabr: Dynamic attribute-based reputation scoring for malicious ip address detection. In *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 64–69. IEEE, 2018.